A Photometric Approach for Estimating Normals and Tangents

Michael Holroyd
Unversity of Virginia

Jason Lawrence
University of Virginia

Greg Humphreys
University of Virginia

Todd Zickler
Harvard University

Abstract

This paper presents a technique for acquiring the shape of real-world objects with complex isotropic and anisotropic reflectance. Our method estimates the local normal and tangent vectors at each pixel in a reference view from a sequence of images taken under varying point lighting. We show that for many real-world materials and a restricted set of light positions, the 2D slice of the BRDF obtained by fixing the local view direction is symmetric under reflections of the halfway vector across the normal-tangent and normal-binormal planes. Based on this analysis, we develop an optimization that estimates the local surface frame by identifying these planes of symmetry in the measured BRDF. As with other photometric methods, a key benefit of our approach is that the input is easy to acquire and is less sensitive to calibration errors than stereo or multi-view techniques. Unlike prior work, our approach allows estimating the surface tangent in the case of anisotropic reflectance. We confirm the accuracy and reliability of our approach with analytic and measured data, present several normal and tangent fields acquired with our technique, and demonstrate applications to appearance editing.

Keywords: Photometric Stereo, BRDF, Symmetry, Anisotropic, Normal Map, Tangent Map

1 Introduction

Measuring the shape and appearance of physical objects remains one of the most challenging problems at the boundary of computer vision and computer graphics. Photometric methods for estimating shape have proved useful for this task because they directly measure surface normals, which are critical for rendering applications or recovering reflectance. While there has been significant progress toward developing "general-purpose" techniques, existing approaches still rely on strong assumptions about the surface properties of the target object that are not always true in practice.

This paper presents a new photometric technique for recovering the shape of opaque objects. As with existing approaches, the input to our algorithm consists of a sequence of images taken under varying point lighting, and the output includes an estimate of the surface normal for each pixel in a reference view. Unlike previous techniques, however, our approach does not require estimating a model of the surface reflectance and performs well for diffuse, glossy, shiny, metallic, dielectric, isotropic, anisotropic materials and for surfaces with homogeneous or spatially-varying reflectance.

Our strategy is to identify symmetries in the 2D slice of the Bidirectional Reflectance Distribution Function (BRDF) [Nicodemus et al. 1977] obtained by fixing the local view direction. In particular, we consider the reflective symmetry of the halfway vector (the bisector of the local view and light directions) across the normal-tangent and normal-binormal planes. Although these symmetries are not guaranteed to be satisfied exactly for all physically-valid BRDFs, we show that when light positions are properly restricted during acquisition, one can reliably estimate normals for a wide range of BRDFs by localizing their planes of symmetry.

In addition to providing a means for recovering surface normals, our symmetry-based approach has the advantage of delivering per-pixel *tangent vectors* as well. For anisotropic surfaces, the tangent field rivals the normal field in its importance for reflectometry and rendering, but to date, its recovery has received very little attention.

This paper presents empirical data for analytic and measured BRDFs demonstrating that our approach produces normals and tangents that are accurate to within a few degrees. We also show normal and tangent fields measured from real-world objects and demonstrate appearance editing applications.

2 Previous Work

Techniques for measuring the shape of real-world objects can be broadly classified as relying either on geometric cues (e.g., identifying a common point seen from different lines of sight) or photometric cues (e.g., observing $(\mathbf{n} \cdot \mathbf{l})$ at multiple light positions). Because our approach relies on a photometric cue (symmetries in the surface BRDF) we review existing photometric methods in this section. To the best of our knowledge, ours is the first approach that allows direct estimation of both the surface orientation and the principal directions of light scattering for anisotropic materials.

Classical photometric stereo technique [Woodham 1980] estimates the surface normal and albedo of a perfectly diffuse (Lambertian) object as seen from each pixel in a fixed view using a set of input images under varying point lighting. Although real-world materials are never perfectly diffuse, the simplicity and stability of the resulting linear optimization are undeniable, and this simple technique can be made more robust using outlier rejection. Additionally, color-based analysis [Mallick et al. 2005] and/or intensity-based analysis [Coleman and Jain 1982] can be used to isolate the diffuse component of more complex reflections provided that a significant diffuse component exists.

Numerous improvements to this basic idea have been proposed, most of which extend to non-Lambertian materials by fitting measurements to low-dimensional BRDF models. Tagare et al. [1991] extended the Lambertian case to a simplified multi-lobe reflectance model. Georghiades [2003] estimates surface orientation and reflectance (as well as light source positions) by assuming the surface can be represented by a simplified Torrance-Sparrow BRDF [1967] with a constant specular component over the surface. Goldman et al. [2005] allow more variation over the surface by simultaneously estimating a set of basis materials modeled as isotropic Ward BRDFs [1992]. A similar method was recently proposed by Chung et al. [2008] which focuses on materials with wide specular lobes. While these approaches apply to a wider class of materials, their

---

$(\mathbf{a} \cdot \mathbf{b})$	scalar dot product of vectors $\mathbf{a}$ and $\mathbf{b}$
$\mathbf{n}$	surface normal
$\mathbf{t}$	surface tangent
$\mathbf{b}$	surface binormal $\mathbf{b} = \mathbf{n} \times \mathbf{t}$
$\mathbf{v} \sim (\theta_v, \phi_v)$	view vector
$\mathbf{l} \sim (\theta_l, \phi_l)$	light vector
$\mathbf{h} \sim (\theta_h, \phi_h)$	halfway vector $\mathbf{h} = (\mathbf{l} + \mathbf{v})/\lVert\mathbf{l} + \mathbf{v}\rVert$
$f_r(\mathbf{v}, \mathbf{l})$	surface BRDF where $\mathbf{v}, \mathbf{l}$ are expressed in the local coordinate system
$\beta(\theta_h, \phi_h)$	2D slice of the BRDF for a fixed view parameterized with respect to the halfway vector

Table 1: *Notation used in this paper. All vectors are normalized.*

dependence on parametric models limits their generality. Natural materials can differ substantially from analytic models [Ngan et al. 2005; Stark et al. 2005], resulting in errors in the deduced shape that can be difficult to characterize. In addition, a single object may exhibit different reflectance at different regions which cannot be expressed by a small linear basis defined over one parametric model. Our approach does not impose any parametric model on the measured reflectance data.

A few techniques also avoid relying on parametric models. Hertzmann and Seitz [2003] place reference objects of known shape and material in the scene to allow modeling a target object as a spatially-varying mixture of these materials. Unfortunately, such reference objects are not practical to create for many scenes. Another example is the work of Alldrin et al. [2008], which represents isotropic reflectance using a linear basis of general non-parametric bivariate functions. Although such an approach holds promise, jointly estimating the component materials and per-pixel surface orientation leads to a fragile optimization that must integrate reflectance across the entire object. On the other hand, our approach processes each pixel independently and can seamlessly handle anisotropic effects.

For shiny objects, a number of methods have been designed to measure surface orientation under the assumption that the point of maximum reflectance occurs when the halfway vector is coincident to the normal [Wang and Dana 2006; Chen et al. 2006; Ma et al. 2007; Francken et al. 2008]. While these systems have generated impressive results for a notoriously difficult class of objects, they are not without limitations. First, observing the specular highlight at every pixel requires a very high sampling rate (or a more complex acquisition system [Nayar et al. 1990; Ma et al. 2007]) and can be physically impossible in concave regions. In contrast, our approach of identifying planes of symmetry is more robust to noisy and sparse data; similar results have been reported in the area of 3D shape analysis where points of symmetry have been shown to be more stable than other statistics [Podolak et al. 2006]. Second, the assumption that the maximum reflectance occurs at $(\mathbf{n} \cdot \mathbf{h}) = 1$ is only valid for mirror-like surfaces. Glossy surfaces exhibit off-specular peaks [Torrance and Sparrow 1967] which bias this estimate. Francken et al. [2008] showed that these methods give systematic errors of up to 5 degrees for glossy materials.

More similar to our approach are methods that exploit symmetries present in real-world BRDFs. Zickler et al. [2002] exploit Helmholtz reciprocity by capturing images in which the position of the camera and light source are interchanged. This enables the recovery of both depth and surface normals independent of reflectance but requires finding corresponding points in images taken from different viewpoints. In contrast, our single-view approach permits a simpler acquisition and allows the direct recovery of surface normals at the resolution of the camera. Isotropy is another form of symmetry that has been exploited for photometric stereo by Alldrin and Kriegman [2007]. For isotropic BRDFs, they show that one component of the surface normal can be recovered by exploiting the reflective symmetry across the view-normal plane. In contrast, the symmetries we consider apply to both isotropic and anisotropic materials and allow measuring the complete surface normal and tangent directions.

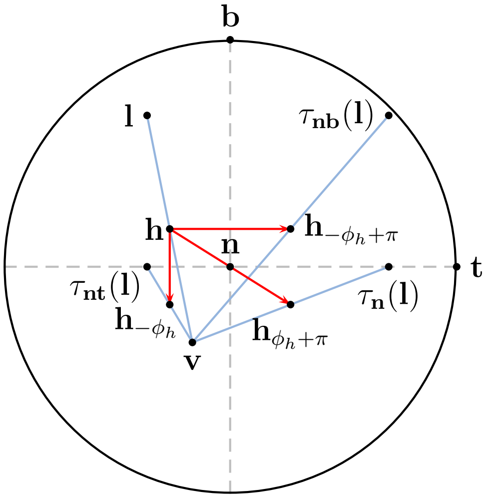

Figure 1: *Orthographic projection of the hemisphere surrounding the surface normal $\mathbf{n}$ annotated with the tangent $\mathbf{t}$ and binormal $\mathbf{b}$ along with the three types of symmetry explored in this paper. For a fixed view $\mathbf{v}$, we expect the value of the BRDF at light position $\mathbf{l}$ to be equal to that at light positions corresponding to reflections of the halfway vector $\mathbf{h}$ across the $\mathbf{nt}$ plane, the $\mathbf{nb}$ plane, and the surface normal.*

3 Theoretical Framework

Our goal is to estimate per-pixel normal and tangent vectors $(\mathbf{n}, \mathbf{t})$ from a sequence of images recorded at different light positions. For a fixed view vector $\mathbf{v}$, let $\beta_{\mathbf{v}}(\theta_h, \phi_h)$ denote the 2D slice of the BRDF sampled at each pixel, parameterized by the elevation and azimuthal angles of the halfway vector expressed in the local coordinate system defined by $(\mathbf{n}, \mathbf{t})$. As illustrated in Figure 1, we explore the following symmetries in these functions:

- Reflection of the halfway vector across the $\mathbf{nt}$ plane:
  $\beta_{\mathbf{v}}(\theta_h, \phi_h) = \beta_{\mathbf{v}}(\theta_h, -\phi_h)$
- Reflection of the halfway vector across the $\mathbf{nb}$ plane:
  $\beta_{\mathbf{v}}(\theta_h, \phi_h) = \beta_{\mathbf{v}}(\theta_h, -\phi_h + \pi)$
- Together, these two symmetries imply reflection across the surface normal:
  $\beta_{\mathbf{v}}(\theta_h, \phi_h) = \beta_{\mathbf{v}}(\theta_h, \phi_h + \pi)$

These symmetries are different from those typically associated with BRDFs such as isotropy, which implies rotation of the light and view around the surface normal; we are considering BRDF values at light/view positions corresponding to reflections of *their associated halfway vectors*. Although such symmetries are not necessary for a physically valid BRDF, we will show that they are present for a wide range of materials and under certain light/view configurations.

Although many simple analytic BRDFs (e.g. Blinn-Phong [1977]) exhibit these symmetries perfectly, it is more informative to consider physically-based models that apply to a much wider class of materials. In particular, microfacet models treat the surface as a collection of microscopic Fresnel mirrors and have been shown to accurately represent a wide variety of both isotropic and anisotropic materials [Cook and Torrance 1982; Ashikhmin et al. 2000; Ngan et al. 2005]. The specular component of a microfacet-based BRDF has the form [Torrance and Sparrow 1967; Ashikhmin et al. 2000]

$$f_r(\mathbf{v}, \mathbf{l}) \propto p(\mathbf{h}) F(\mathbf{l} \cdot \mathbf{h}) \frac{S(\mathbf{v}, \mathbf{l}, \mathbf{h})}{(\mathbf{v} \cdot \mathbf{n})(\mathbf{l} \cdot \mathbf{n})}, \tag{1}$$

where $p(\mathbf{h})$ captures the distribution of microfacet orientations, F is the Fresnel term, and $S(\mathbf{v}, \mathbf{l}, \mathbf{h})$ models shadowing and masking effects at the micro scale. The distribution $p(\mathbf{h})$ is commonly modeled with an analytic function such as an elliptical Gaussian [Ward 1992] or the Beckmann distribution [Cook and Torrance 1982]. Alternatively, Ngan et al. [2005] have shown that it can be fit to measured data using Ashikhmin et al.'s [2000] approximation to Equation 1.

In the context of the symmetries described above, one can make two observations regarding the microfacet model: 1) the distribution term $p(\mathbf{h})$ has a dominating influence on appearance for a broad range of view and light directions; and 2) it generally exhibits the symmetries that we seek to exploit. The latter observation comes from existing analytical and measured distributions, almost all of which satisfy the symmetry conditions. While exceptions do exist (see Section 6), we assume they are rare. The first observation is based on both physical and empirical evidence. Empirically, while the right-most term in Equation 1 can be quite complex, it is typically smooth and can be simplified tremendously without significant loss of visual accuracy [Ashikhmin et al. 2000; Ngan et al. 2005]. Physically, the other term to consider—the Fresnel term—is approximately proportional to $(1 - (\mathbf{l} \cdot \mathbf{h}))^5$ [Schlick 1994], and is substantial only at grazing angles. By restricting the set of light positions to a limited cone of directions about the view vector, the influence of this term can be made small as well.

In the following, we design our acquisition system to include the cone of light directions for which $\cos^{-1}(\mathbf{l} \cdot \mathbf{h}) < \theta_{d_{\max}}$ (the notation θ_d is due to Rusinkiewicz [1998]; it refers to the elevation angle of the *difference* vector). Since θ_d depends only on the light and view directions, this bound can be achieved by simply restricting the light source positions during acquisition without any prior knowledge of the surface normal. By bounding the light directions in this way, we ensure the dominance of the microfacet distribution. Our results show that it is then possible to robustly identify the planes of reflective symmetry from the partial slice of the BRDF for a wide range of surface orientations.

3.1 A Measure of BRDF Symmetry

In this section, we define a precise measure of the symmetry previously discussed. For a fixed view, let $\mathbf{l}$ and $\mathbf{l}'$ be two light positions with corresponding halfway vectors at symmetric positions across some plane. The measured intensities at these positions will be $I_{\mathbf{l}} \propto f_r(\mathbf{v}, \mathbf{l})(\mathbf{n} \cdot \mathbf{l})$ and $I_{\mathbf{l}'} \propto f_r(\mathbf{v}, \mathbf{l}')(\mathbf{n} \cdot \mathbf{l}')$. We ignore the intensity of the light which is the same in all measurements and therefore does not affect our analysis. Note that even if the BRDF is equal at these positions, the product of the BRDF and the cosine term need not be. This can easily be corrected by cross-multiplying the cosine terms. In other words, if the BRDF is indeed symmetric across the chosen plane then $(\mathbf{n} \cdot \mathbf{l}')I_{\mathbf{l}} = (\mathbf{n} \cdot \mathbf{l})I_{\mathbf{l}'}$.

Now let $\tau(\mathbf{l})$ be an arbitrary transformation on lighting positions. For some hypothesized normal $\mathbf{n}$ and tangent $\mathbf{t}$, we define the *symmetry distance* of the BRDF under this transformation as:

$$\mathrm{SD}_\tau(\mathbf{n}, \mathbf{t}) = \frac{\int_{\Omega_\tau} \|(\mathbf{n} \cdot \tau(\mathbf{l}))I_{\mathbf{l}} - (\mathbf{n} \cdot \mathbf{l})I_{\tau(\mathbf{l})}\|^2 \, d\omega_l}{\int_{\Omega_\tau} \|(\mathbf{n} \cdot \tau(\mathbf{l}))I_{\mathbf{l}}\|^2 \, d\omega_l},$$

where the domain of integration Ω_τ is constrained to include only light positions $\mathbf{l}$ and $\tau(\mathbf{l})$ that lie in the upper hemisphere with respect to $\mathbf{n}$ and correspond to halfway vectors between 0 and $\theta_{d_{\max}}$. As shown in Figure 1, we define $\tau_{\mathbf{nt}}(\mathbf{l})$ to reflect the halfway vector associated with $\mathbf{l}$ across the $\mathbf{nt}$ plane, $\tau_{\mathbf{nb(l)}}$ to reflect the halfway vector across the $\mathbf{nb}$ plane, and $\tau_{\mathbf{n}}$ to reflect the halfway vector across the normal. Note that the transformation $\tau_{\mathbf{n}}$ is independent of the chosen tangent direction. Finally, we define a combined symmetry measure as

$$\mathrm{SD}(\mathbf{n}, \mathbf{t}) = \mathrm{SD}_{\tau_{\mathbf{nt}}}(\mathbf{n}, \mathbf{t}) + \mathrm{SD}_{\tau_{\mathbf{nb}}}(\mathbf{n}, \mathbf{t}) + \mathrm{SD}_{\tau_{\mathbf{n}}}(\mathbf{n}). \tag{2}$$

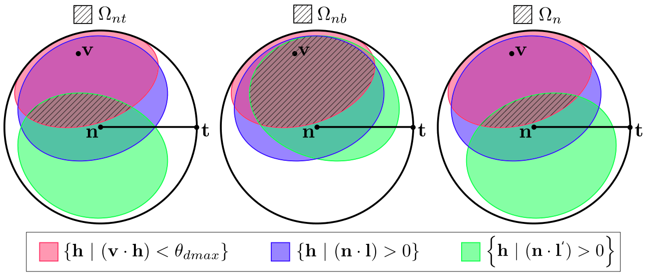

Figure 2: *Integration regions for the three different transformations we consider shown over halfway vectors. These are defined as the intersection of the sets of halfway vectors for which* $(\mathbf{n} \cdot \mathbf{l}) > 0$ *and* $(\mathbf{n} \cdot \mathbf{l}') > 0$ *that also have* $\cos^{-1}(\mathbf{v} \cdot \mathbf{h})$ *within* $\theta_{d_{\max}}$. *These regions can form complex shapes and can even become empty for normals that lie beyond* $\theta_{d_{\max}}$.

When $\mathbf{n}$ and $\mathbf{t}$ coincide with the true surface frame, we expect SD to be minimized.

As illustrated in Figure 2, the integration regions $\Omega_{\mathbf{nt}}$, $\Omega_{\mathbf{nb}}$ and $\Omega_{\mathbf{n}}$ typically form complex shapes and are even empty for certain combinations of $\mathbf{n}$, $\mathbf{t}$, and $\theta_{d_{\max}}$. Therefore, care must be taken when evaluating these integrals as discussed in the following section.

3.2 Validation

We performed numerous simulations designed to verify our key assumption that the symmetry distance is minimized at the correct surface frame and to evaluate the effect of $\theta_{d_{\max}}$. This was done using a standard Monte Carlo algorithm to estimate Equation 2 at a dense set of normal and tangent vectors. We compared the normal that maximized symmetry to the ground truth normal for a number of analytic and measured BRDFs as well as for different values of $\theta_{d_{\max}}$ and surface orientations.

We first considered anisotropic materials using the elliptical Gaussian Ward model [Ward 1992] along with three of the measured anisotropic samples acquired by Ngan et al. [2005]. For the measured data, we followed their reported procedure for fitting the tabulated microfacet distribution, Fresnel parameter, and shadowing and masking terms of the Ashikhmin model [2000], and these fits were used for our experiments. For our simulations, we fixed the view vector to be the z-axis, chose a normal vector within the zx plane, and fixed the tangent at an angle $\phi_t = 25°$ off this plane. We estimated the normal and tangent vectors that minimize Equation 2 and compared these to the correct surface frame.

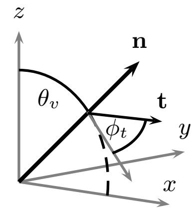

Figure 3 reports the errors in these estimates for each material over a range of values of $\theta_{d_{\max}}$ (varying along the columns in each

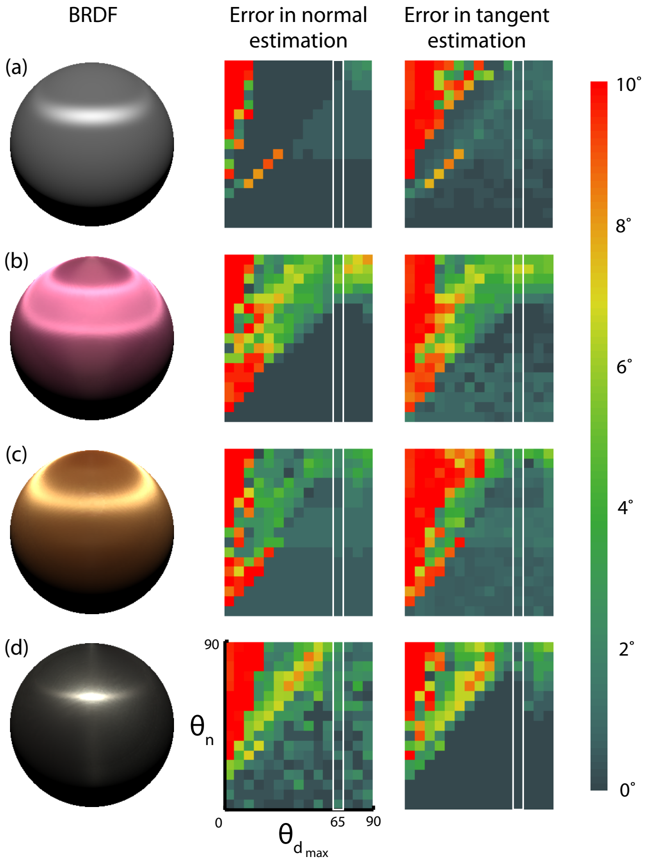

Figure 3: *Error analysis using the Ward anisotropic BRDF and three measured samples [Ngan et al. 2005]. The errors in the estimated normal and tangent vectors are reported in degrees over a range of reflectance angles θ_n (the angle formed by the view and normal) which vary across the rows and values of $\theta_{d_{\max}}$ which vary across the columns. The materials include (a) an elliptical Ward with parameters $k_d = 0.5$, $k_s = 0.5$, $\alpha_x = 0.1$, and $\alpha_y = 0.5$.; (b) measured purple satin; (c) measured yellow satin; and (d) measured brushed aluminum.*

false-color error visualization) and a range of surface orientations (varying along the rows). These visualizations show two important trends. First, our ability to identify the correct normal and tangent diminishes as $\theta_{d_{\max}}$ diminishes. For very restrictive sets of light positions, most candidate normals result in a very small region of integration (see Figure 2), making it difficult to reliably identify the minimum. Therefore, we disregard estimates at pixels where we identify any of these domains to be empty. This limits the range of recovered normals to those less than $\theta_{d_{\max}}$ since $\Omega_{\mathbf{n}}$ is always empty beyond this point.

The second important trend is that the accuracy of our approach steadily increases as $\theta_{d_{\max}}$ approaches 70° degrees and then falls off for larger values. We attribute this to the presence of Fresnel effects and the influence of shadowing and masking components in these regions of the BRDF as previously discussed. Based on this analysis, we chose $\theta_{d_{\max}} = 65^\circ$ as the "sweet spot" (indicated by the white boundary in Figure 3) and used this for all the datasets we captured. Note that for these materials our approach typically identified the correct normal and tangent to within a single degree and only at extreme orientations did we observe errors of up to four degrees. These results are representative of those we observed at different tangent vectors and we chose to present those for $\phi_t = 25^\circ$ only for the sake of clarity.

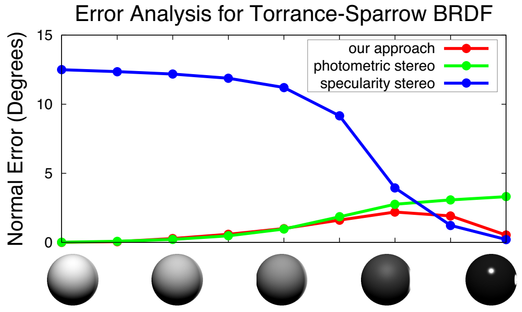

Figure 4: *Average error computed for Lambertian photometric stereo with outlier rejection, our approach, and a method that associates the direction of maximum reflectance with the normal for different materials ranging from perfectly diffuse to dark and shiny modeled with the Torrance-Sparrow BRDF [1967].*

We also experimented with isotropic materials using the Torrance-Sparrow BRDF [Torrance and Sparrow 1967] with parameters ranging from perfectly diffuse to dark and shiny. The results are shown in Figure 4, where the average error was computed over normal orientations ranging from 0° to 60° and for $\theta_{d_{\max}} = 65^\circ$ (recall that the value of $\theta_{d_{\max}}$ places a restriction on the range of normals that our method can reliably estimate). As a baseline comparison, we also report the average error for Lambertian photometric stereo [Woodham 1980] and specularity stereo [Wang and Dana 2006; Chen et al. 2006; Ma et al. 2007; Francken et al. 2008], which assumes the maximum reflectance occurs when the halfway vector is coincident to the normal. We applied photometric stereo to a simulated dataset consisting of 90×90 light positions sampled uniformly over the sphere and rejected data greater than two standard deviations away from the mean in order to remove areas in shadow or highlight. For the specularity method, we found the maximum of the product of the BRDF and cosine term using a non-linear search; this gives an upper-bound on the performance of these techniques since the angular sampling rate of any practical acquisition device would further reduce accuracy. For all parameter values, our method performs well, and the surface normal is recovered to within a few degrees. Better performance is achieved for very specular surfaces, in which case the distribution term clearly dominates, and for very diffuse surfaces, in which case the $(\mathbf{n} \cdot \mathbf{l})$ terms cancel out only at the correct surface frame. While traditional photometric stereo and specularity-based methods perform well for diffuse and shiny materials respectively, our approach performs well in both cases. We found these trends and absolute errors to be consistent with other analytic models including Cook-Torrance [1982], Ward [1992], and He-Torrance-Sillion-Greenberg [1991].

4 Implementation

Our acquisition setup consists of a QImaging Retiga 4000R digital camera and a tungsten-halogen light source mounted to a computer-controlled spherical gantry [Stanford 2002]. The light positions returned by the gantry are accurate to within one tenth of a degree. We assume distant lighting and an orthographic camera model. Note that a less precise setup would be possible if used in conjunction with scene fiducials that allow estimating the position of a light source [Masselus et al. 2002; Chen et al. 2006; Toler-Franklin et al. 2007].

At each pixel in the reference view we reconstruct the BRDF slice $\beta_{\mathbf{v}}(\theta_h, \phi_h)$ using spherical barycentric interpolation over the measurements [Cabral et al. 1999]. As a pre-process, we compute the Delaunay triangulation [Shewchuk 1996] of the halfway vec-

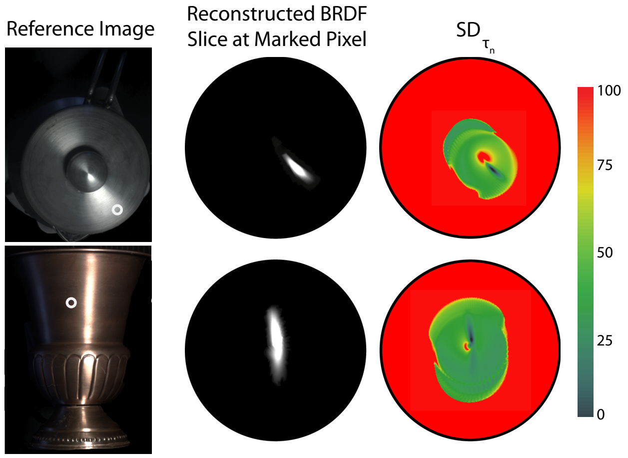

Figure 5: *Typical reconstructions of the BRDF slice shown for two anisotropic datasets along with a visualization of the symmetry distance at a dense set of normal vectors. The BRDF slices and the symmetry values are functions defined over the hemisphere which are shown here using a parabolic projection onto the plane [Heidrich and Seidel 1999]. Symmetry distances are normalized to occupy the range* $[0, 100]$.

tors corresponding to measurement locations projected onto the unit disc. Figure 5 (middle) visualizes the reconstructed BRDF slices at a single pixel in two anisotropic datasets. These slices are functions defined over the hemisphere shown here using a parabolic projection [Heidrich and Seidel 1999] onto the plane.

In practice, the accuracy of this reconstruction is dependent on the frequency content of the underlying BRDF, the density and pattern of light source positions, and the interpolation method. Each of the four datasets we acquired consists of approximately $1,500$ high dynamic range images recorded at light positions sampled uniformly over the cone of directions surrounding the view for which $\theta_d < 65^\circ$. This acquisition process requires approximately 45 minutes. The impact of lower sampling densities is examined further in section 6.

Given $\beta_\mathbf{v}(\theta_h, \phi_h)$, we compute the normal and tangent vectors in two steps. First, we find the normal vector that minimizes $\mathrm{SD}_{\tau_\mathbf{n}}$. Recall that this distance is independent of the choice of tangent direction even for anisotropic materials. This involves performing a non-linear optimization over a 2D domain for which we use the Nelder-Mead simplex algorithm [1965]. We initialize this search at the direction where we observed the maximum reflectance. Figure 5 (right) visualizes our estimate of $\mathrm{SD}_{\tau_\mathbf{n}}$ over a dense range of possible normal directions. The distance at each candidate normal $\mathbf{n}$ is computed as a discrete sum over the sampled light positions. For each light position, we reflect its halfway vector across $\mathbf{n}$ and evaluate our reconstructed $\beta_\mathbf{v}(\theta_h, \phi_h)$ at that position. Reflected positions that fall outside $\Omega_\mathbf{n}$ are discounted. In cases where all of the light positions are reflected to areas without data, we assign an arbitrarily large error value (visible as the constant red area along the boundary in Figure 5). These visualizations clearly illustrate the stability of our optimization; the correct minimum is clearly defined and the error space is largely free of poor local minima. These properties result in a robust search that converges quickly, requiring only 60 evaluations of the error function on average.

In the second step, we hold the normal vector fixed and search for the tangent vector that minimizes $\mathrm{SD}_{\tau_\mathbf{nt}} + \mathrm{SD}_{\tau_\mathbf{nb}}$. Again, we use the Nelder-Mead algorithm to perform this 1D search. Because there exists a natural ambiguity between the binormal and tangent, we use a simple heuristic that assigns the tangent to the direction along which the specular highlight is widest.

We also found it was important to assign a confidence to each pixel based on the residual error and the magnitude of the denominator in Equation 2 at the computed frame. We consistently observed low confidence values in areas where the normal either lies beyond 60° from the view or where the reflectance data is corrupted by interreflections or large cast shadows. We cull any estimate whose confidence is below a user-set threshold, which is determined manually for each dataset. Since each pixel may be processed independently, we distributed this optimization over a cluster of forty-two machines with Dual 1.6 GHz Opteron252 CPUs and 2GB of memory. Each dataset consists of $1,500$ images with 1024×1024 resolution and requires 2.3 GB of disk space. The optimization required approximately 10 minutes to complete or the equivalent of 7 hours on a single machine.

5 Results

Figure 6 shows the normal fields we estimated for two isotropic datasets. The first dataset is of a vase with an intricate spatially-varying design and a shiny surface finish. The second is of a small wooden figurine in the shape of a frog also with a spatially-varying appearance, but a more matte finish. We have masked out regions with confidence values below our chosen threshold (shown as a dotted line in the reference image). The few errors we observe are primarily caused by scene interreflections (e.g., near the top of the vase in Figure 6). Despite these errors, our algorithm is able to capture the overall shape of these objects even though they exhibit significantly different material properties. Furthermore, our technique recovers fine geometric details such as the embossed pattern on the vase and small indentations in the frog figurine as seen in the cutouts.

Figure 7 shows normal and tangent fields computed from two objects with anisotropic reflectance. The first dataset is of the top of a french press coffeemaker which shows aluminum brushed in a circular pattern. The second dataset is a bronze vase which has been brushed in a constant direction around its axis of symmetry. The tangent fields computed with our algorithm correctly capture these patterns, although there are regions where the material is not brushed and therefore the tangent direction is undefined (e.g., near the ridges in the middle of the vase). Additionally, errors in the normals often lead to errors in the tangents as seen near the bottom of the bronze vase.

5.1 Application: Appearance Editing

Our method provides separate estimates of the 3D shape of an object and its surface reflectance, allowing for independent manipulation of either component. Figure 8 shows an example in which the surface reflectance is edited to enhance surface gloss. In this example, we enable control over the reflectance by independently fitting an analytic Ward BRDF to the reflectance measurements available at each surface position. A comparison of this fit to a reference image suggests that the Ward model provides a reasonable description of this surface's reflectance, at least for this particular slice of the BRDF domain. The right of this figure shows an editing example in which we have scaled the specular and diffuse components at each pixel to give the finish a glossier appearance.

Since we use a parametric BRDF model for editing, this approach is conceptually similar to what could be achieved using the method of Goldman et al. [2005]. An important difference, however, is that our approach recovers geometry in a manner that is separate

Figure 6: *Normals computed from two objects with complex spatially-varying isotropic reflectance. Normals are visualized according to* $(r, g, b) = (\frac{n_x+1}{2}, \frac{n_y+1}{2}, n_z)$*. The dashed line in each reference image indicates the region shown at right.*

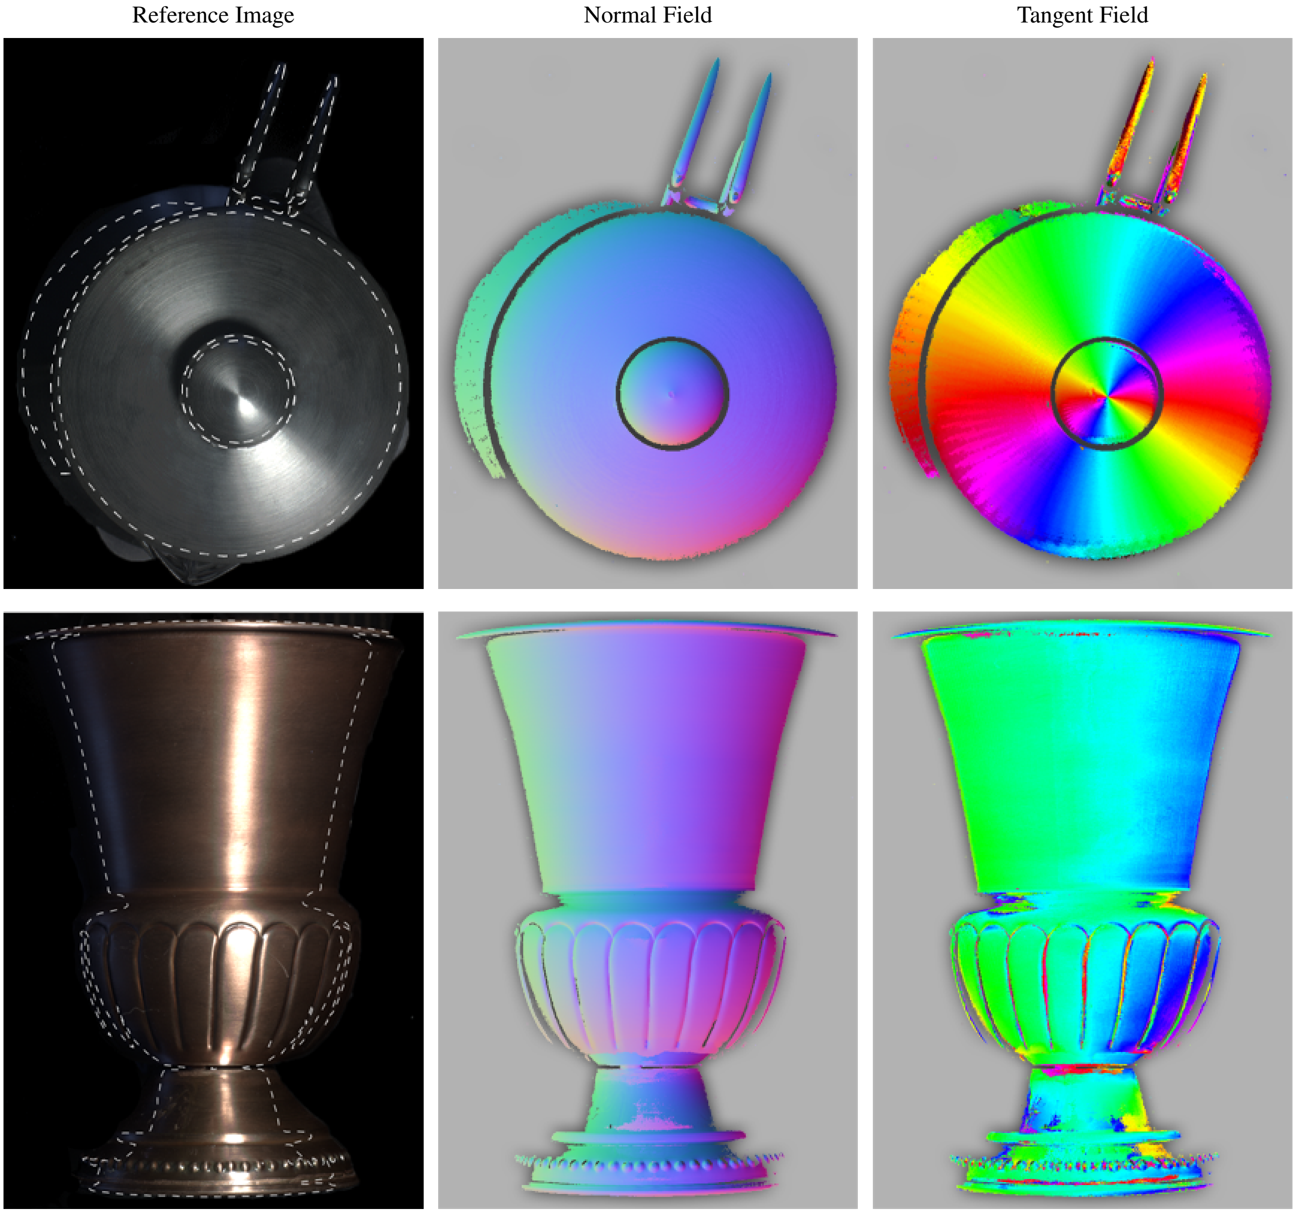

Figure 7: *Normals and tangents computed from two objects with complex anisotropic reflectance. Tangent directions are visualized according to* $(h, s, v) = (2\phi, 1, 1)$ *where the angle* ϕ *is defined with respect to the x-axis.*

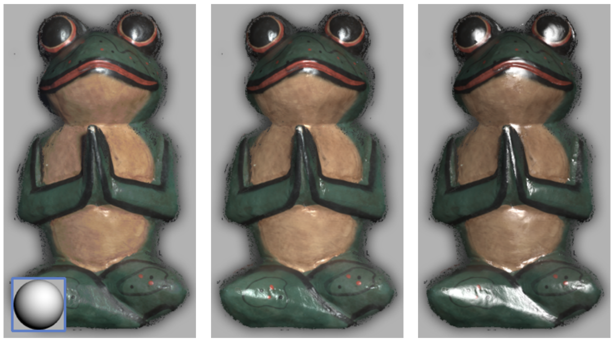

Reference Image Ward Fits Edited Fits

Figure 8: *Edited surface reflectance combined with measured geometry.* **Left:** *original HDR image not included in training set.* **Middle:** *reconstruction from fits of the Ward BRDF model computed at each pixel.* **Right:** *result of editing the specular component to give the surface a glossier appearance.*

Original Edited Tangents

Figure 9: *Edited tangent field combined with measured reflectance.* **Left:** *original HDR image not included in training set.* **Right:** *image generated by editing the tangent field to display the SIGGRAPH logo.*

from reflectance. This provides the ability to choose editing tools that are most appropriate for the reflectance of a particular object. These tools can be based on a parametric model appropriate to the object—as we have demonstrated here—or on data-driven representations similar to those described by Lawrence et al. [2006].

It is also possible to manipulate the underlying normal and tangent fields of a surface while keeping the original reflectance intact. Figure 9 shows a relighting result that uses an edited version of the bronze vase's tangent field to display the SIGGRAPH logo. This edit was accomplished by rotating each tangent vector within the edited region by 90° about its corresponding surface normal; Fisher et al. [2007] have recently introduced tools that enable more sophisticated manipulations of tangent fields.

6 Discussion and Future Work

While we have shown that our strategy of associating planes of symmetry in the BRDF with the correct surface frame is valid for many different materials, we did observe cases where this approach fails. A good example is the red velvet sample acquired by Ngan et al. [2005]. The deduced microfacet distribution does not exhibit the expected symmetries and we calculated that our technique would produce errors of up to ten degrees. Intuitively, materials with microgeometries symmetric around the surface normal (grooves in brushed metal, threads in satin) satisfy our underlying assumption, whereas materials with asymmetric microgeometries such as velvet (composed of tiny cylinders systematically brushed to a fixed angle off the normal) do not. Retroreflective materials would also present challenges for our approach. These BRDFs exhibit backscattering lobes that are substantial when the light and view directions are close to one another. However, we expect it would be possible to handle these cases by placing a lower-bound on $\theta_{d_{\max}}$ during acquisition to avoid sampling these regions.

Accurately identifying the planes of symmetry in $\beta_{\mathbf{v}}(\theta_h, \phi_h)$ depends strongly on our ability to reconstruct this function from our sampled data. Figure 10 shows the effect on the resulting normal field of using a decreasing number of light source positions. As expected, with fewer source positions the quality of these reconstructions decrease and the resulting normal field exhibits a faceted appearance as the optimization more often returns a poor local minimum. Using a more sophisticated interpolation technique such as radial basis functions would result in better reconstructions and is an interesting direction of future work.

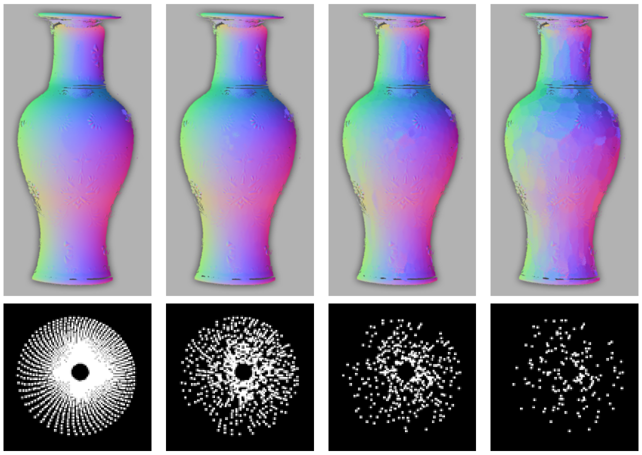

1,512 743 380 172

Figure 10: *Normal maps of a porcelain vase computed using different numbers of light positions. The images along the bottom visualize the source positions showing the corresponding half-way vector projected onto the x-y plane of a view-centered coordinate system.*

Another limitation of our technique is that it does not account for scene interreflections or subsurface scattering. Figure 11 visualizes the normal and tangent map of a silk necktie computed using our approach. Note that this scene exhibits strong discontinuities in the tangent field and many depth boundaries. The majority of errors are caused by interreflections which occur near the overlapping regions. Future research might consider combining our approach with techniques for isolating light that is directly scattered from the object surface such as that presented by Nayar et al. [2006]. We should also note that in some places our heuristic for choosing between the tangent and the binormal is unreliable (e.g., the abrupt blue-to-yellow transition in the tangent map visualization in the bullseye pattern near the bottom left). This is due to the fact that the highlights for this material are relatively wide and sparsely sampled near grazing angles. Despite these few inconsistencies, however, the majority of this challenging scene is accurately captured by our method.

It is also worth emphasizing that there are a number of other photometric stereo methods that could be used in conjunction with our approach. For example, in the isotropic case, symmetries along

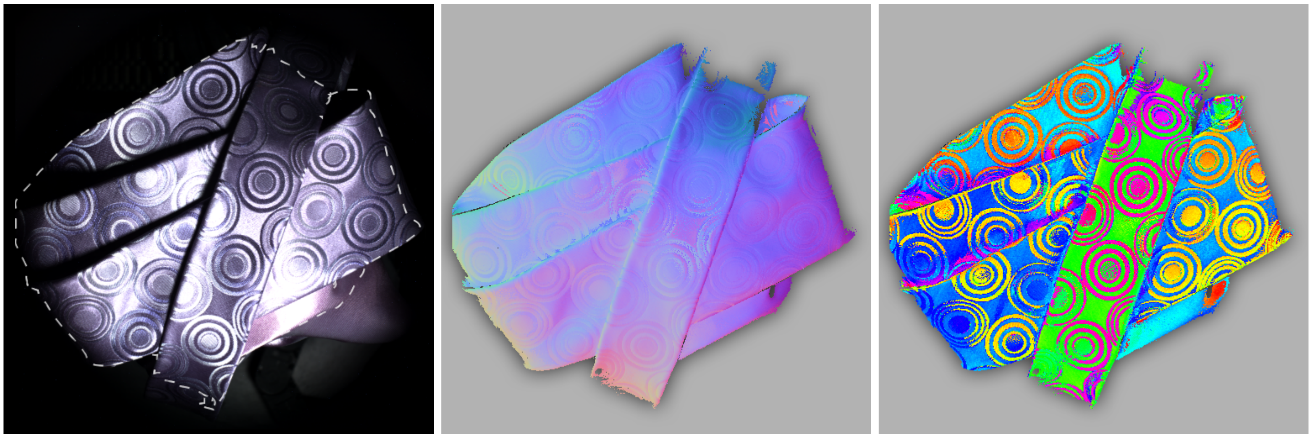
Reference Image Normal Field Tangent Field

Figure 11: *A folded necktie with complex reflectance, depth boundaries and discontinuities in the tangent field presents a particularly challenging case.*

intensity profiles from view-centered circles of light directions provide an additional constraint on the surface normal [Alldrin and Kriegman 2007]. Rather than explore the benefits of combining complimentary approaches with our symmetry-based technique, we instead explored what can be recovered using these symmetries alone. This allows us to consider isotropic, anisotropic, diffuse, specular and hybrid surfaces in a unified manner.

Finally, there are many applications for which partial and even imprecise surface geometry is sufficient. A good example is the system recently proposed by Toler-Franklin et al. [2007] for generating stylized illustrations of objects using color images augmented with per-pixel surface normals obtained with Lambertian photometric stereo. The applications described in Section 5 are other examples. We believe that the simplicity of the acquisition required for our technique along with its applicability to a wide range of materials will have broad applicability in similar systems.

7 Conclusion

This paper introduced a new photometric approach for estimating the surface normals and, for anisotropic materials, the tangent vectors over an object's surface as seen from a fixed viewpoint. The input to our algorithm is a set of images taken under variable point lighting constrained to lie within a cone of directions centered at the view. We show that the size of this cone can be chosen to allow reliable detection of the planes of reflective symmetry in the microfacet distribution of the measured BRDF. We demonstrated an optimization that computes the surface frame at each pixel by identifying these planes of symmetry and showed that our approach is accurate to within a few degrees over a range of analytic and measured BRDF data. Finally, we presented normal and tangent fields computed with our technique from four real-world objects that encompassed a range of complex isotropic, anisotropic, and spatially-varying reflectance.

Acknowledgments

The authors thank Jiajun Zhu for discussions and help with data capture. This work was supported by an NSF CAREER award CCF-0747220 and NSF grant CCF-0811493. Jason Lawrence acknowledges an NVIDIA Professor Partnership award and Todd Zickler acknowledges a fellowship from the Sloan Foundation.

References

ALLDRIN, N. G., AND KRIEGMAN, D. J. 2007. Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 1–8.

ALLDRIN, N., ZICKLER, T., AND KRIEGMAN, D. 2008. Photometric stereo with non-parametric and spatially-varying reflectance. In *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR)*.

ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH*, 65–74.

BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. In *Computer Graphics (Proceedings of ACM SIGGRAPH)*, 192–198.

CABRAL, B., OLANO, M., AND NEMEC, P. 1999. Reflection space image based rendering. In *Proceedings of ACM SIGGRAPH*, 165–170.

CHEN, T. B., GOESELE, M., AND SEIDEL, H. P. 2006. Mesostructure from specularity. In *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR)*, 1825–1832.

CHUNG, H.-S., AND JIA, J. 2008. Efficient photometric stereo on glossy surfaces with wide specular lobes. In *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR)*.

COLEMAN, E., AND JAIN, R. 1982. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Computer Vision, Graphics and Image Processing 18*, 4, 309–328.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. In *Computer Graphics (Proceedings of ACM SIGGRAPH)*, 7–24.

FISHER, M., SCHRODER, P., DESBRUN, M., AND HOPPE, H. 2007. Design of tangent vector fields. *ACM Transactions on Graphics (TOG) 26*, 3.

FRANCKEN, Y., CUYPERS, T., MERTENS, T., GIELIS, J., AND BEKAERT, P. 2008. High quality mesostructure acquisition using specularities. In *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR)*.

GEORGHIADES, A. S. 2003. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 816–823.

GOLDMAN, D. B., CURLESS, B., HERTZMANN, A., AND SEITZ, S. M. 2005. Shape and spatially-varying BRDFs from photometric stereo. In *Proceedings of the International Conference on Computer Vision*, 341–348.

HE, X. D., TORRANCE, K. E., SILLION, F., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. In *Computer Graphics (Proceedings of ACM SIGGRAPH)*.

HEIDRICH, W., AND SEIDEL, H.-P. 1999. Realistic, hardware-accelerated shading and lighting. In *Proceedings of ACM SIGGRAPH)*, 171–178.

HERTZMANN, A., AND SEITZ, S. 2003. Shape and materials by example: a photometric stereo approach. *IEEE Computer Vision and Pattern Recognition, 2003. 1*, 533–540.

LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHI, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (TOG) 25*, 3, 735–745.

MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of Eurographics Symposium on Rendering*.

MALLICK, S., ZICKLER, T., KRIEGMAN, D., AND BELHUMEUR, P. 2005. Beyond Lambert: Reconstructing specular surfaces using color. *Proc. IEEE Conf. Computer Vision and Pattern Recognition 2*, 619–626.

MASSELUS, V., DUTRÉ, P., AND ANRYS, F. 2002. The free-form light stage. In *Proceedings of the Eurographics Workshop on Rendering*.

NAYAR, S., IKEUCHI, K., AND KANADE, T. 1990. Determining shape and reflectance of hybrid surfaces by photometricsampling. *Robotics and Automation, IEEE Transactions on 6*, 4, 418–431.

NAYAR, S. K., KRISHNAN, G., GROSSBERG, M. D., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics 25*, 3.

NELDER, J., AND MEAD, R. 1965. A simplex method for function minimization. *Computer Journal 7*, 308–311.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*, 117–126.

NICODEMUS, F. E., RICHMOND, J. C., AND HSIA, J. J. 1977. Geometrical considerations and reflectance. *National Bureau of Standards*.

PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Transactions on Graphics 25*, 3.

RUSINKIEWICZ, S. 1998. A new change of variables for efficient BRDF representation. In *Proceedings of the Eurugraphics Rendering Workshop*, 11–22.

SCHLICK, C. 1994. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum 13*, 3, 233–246.

SHEWCHUK, J. R. 1996. Triangle: Engineering a 2D quality mesh generator and delaunay triangulator. In *Applied Computational Geometry: Towards Geometric Engineering*, M. C. Lin and D. Manocha, Eds., vol. 1148 of *Lecture Notes in Computer Science*. Springer-Verlag, may, 203–222.

STANFORD, 2002. Stanford graphics lab spherical gantry (http://graphics.stanford.edu/projects/gantry/).

STARK, M., ARVO, J., AND SMITS, B. 2005. Barycentric parameterizations for isotropic BRDFs. *IEEE Transactions on Visualization and Computer Graphics*, 126–138.

TAGARE, H. D., AND DEFIGUEIREDO, R. J. P. 1991. A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence 13*, 2.

TOLER-FRANKLIN, C., FINKELSTEIN, A., AND RUSINKIEWICZ, S. 2007. Illustration of complex real-world objects using images with normals. In *International Symposium on Non-Photorealistic Animation and Rendering*.

TORRANCE, K. E., AND SPARROW, M. E. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America 57*, 1105–1114.

WANG, J., AND DANA, K. J. 2006. Relief textures from specularities. *IEEE Transactions on Pattern Analysis and Machine Intelligence 28*, 3, 446–457.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of ACM SIGGRAPH)*, 265–272.

WOODHAM, R. J. 1980. Photometric method for determining surface orientation from multiple images. *Optical Engineering 19*, 1, 139–144.

ZICKLER, T., BELHUMEUR, P. N., AND KRIEGMAN, D. J. 2002. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision 49*, 2-3, 215–227.